

STAT 0311: STATISTICAL INFERENCE

Spring 2025

Instructor: Christian Stratton	Time: TR 11:15 – 12:30
Email: cstratton@middlebury.edu	Place: 75 Shannon St 224
Office: Warner 203	Office hours: TBD Also by appointment

Course description: An introduction to the mathematical methods and applications of statistical inference using both classical methods and modern resampling techniques. Topics will include: permutation tests, parametric and nonparametric problems, estimation, efficiency and the Neyman-Pearsons lemma. Classical tests within the normal theory such as F-test, t-test, and chi-square test will also be considered. Methods of linear least squares are used for the study of analysis of variance and regression. There will be some emphasis on applications to other disciplines.

Correspondence: My goal is to maximize my availability for help and discussion throughout the semester. Office hours will be determined via poll during the first week of class, but please feel free to contact me via email at anytime. Additionally, I am happy to meet outside of office hours by appointment.

Meeting format: Class time will generally be used to learn new statistical concepts through a mixture of lecture and in-class activities. Most class periods will feature a short lecture introducing a new concept, followed by an in-class guided activity to be worked on in small groups. You will need to have access to a laptop during class.

Learning objectives: This course serves as an introduction to point estimation, interval estimation, and hypothesis testing and is meant to prepare you for future coursework in statistics. In addition to learning mathematics for common statistical methodology, we will discuss statistical inference and you will learn to think and reason statistically. Through this course, students will:

- Derive estimators, hypothesis tests and confidence intervals for unknown parameters
- Compare and evaluate estimators, hypothesis tests and confidence intervals based on desirable statistical properties
- Demonstrate an understanding of how computer simulation may be used to investigate properties of estimators
- Apply concepts to practical problems and relate them to other coursework and experiences in statistics.

Textbook and materials: There is nothing that need be purchased for this class.

- The website for this course is on Middlebury Canvas. Please check Canvas often for assignments, deadlines, resources, and announcements.
- Students must have access to a laptop with the statistical computing language R, which can be downloaded for free at <https://cran.rstudio.com/>. Additionally, I recommend using RStudio as an integrated development environment (IDE) for interfacing with R. RStudio may be downloaded for free at <https://posit.co/download/rstudio-desktop/>.
 - Laptops with R/RStudio pre-installed are available to borrow from the Davis Family Library, which are a good option for those without access to a laptop or those experiencing short-term issues with your laptop. Please talk to me or the front desk of the Davis Library for more info.

- Lecture notes will be based on *Mathematical Statistics with Applications*, 7th ed. by Wackerly, Mendenhall, and Scheaffer. Please see me if you would like help accessing this text. Additional supplementary texts include:
 - *Statistical inference (2nd edition)* by Casella and Berger.
 - *Mathematical statistics and data analysis (3rd edition)* by John A. Rice

Academic integrity: You are bound by Middlebury College’s honor code, including its policies on plagiarism and cheating. Violation of these rules is ground for failure. To avoid charges of plagiarism, cite all the sources used to complete your assignments/homework, including any peers with whom you collaborated. I encourage you to seek help in understanding the concepts and problems in your assignments from various sources, including peers, instructors, peer tutors, class notes, textbooks, and online sources.

Use of LLM and generative AI: Large language models (LLM) and generative AI, such as [ChatGPT](#), are powerful tools enabled by statistics and data science techniques that may be used to enhance your learning of statistics and coding languages. As such, the use of large language models (LLM) and generative AI, such as ChatGPT, is permitted in this class and may be used on all homework assignments, take-home exams, and projects. However, **you may not copy responses verbatim from these tools, nor may you use these tools to generate complete responses or assignments.** Additionally, if content from generative AI is used on an assignment, **you must provide appropriate citation.** To clarify this policy, examples of acceptable and unacceptable prompts for ChatGPT are provided below.

Acceptable:

- What does it mean for an estimator to be unbiased?
- I am writing code to demonstrate that the statistic $T(X) = \frac{1}{n} \sum_i (X_i - \mu)^2$ is a biased estimated estimate of the population variance. My code doesn’t seem to be working, could you...

Unacceptable:

- Prove that the sample variance is an unbiased estimate of the population variance.
- Write code that simulates...
- Answer the following question: *copy-paste from assignment*

Disclaimer: I am compelled to note that while generative AI can be a powerful tool, it is not infallible and genuinely quite bad at statistical inference. Consider the toy exchange provided at the end of the syllabus, conducted on ChatGPT 4o mini on 2024/09/01. For a context-specific example, see the second exchange (Chat is wrong). It is possible that generative AI will provide you with incorrect information, and it is your responsibility to use generative AI critically. “ChatGPT said so,” is not sufficient justification for an answer, and I am unlikely to be sympathetic to such comments on assignments.

Late policy: I understand that extenuating circumstances may arise. If a circumstance arises that prevents you from completing an assignment on time, please notify me as soon as possible. Otherwise, the late policy depends upon the assignment type:

- **Homeworks** Each assignment will come with a recommended due date, which is the date by which I recommend the assignment is completed to keep up with course material. You may turn an assignment in up to three weeks after the recommended due date with no penalty. During that period, you may resubmit any assignment and I will regrade it with no penalty. No late work will be accepted after those three weeks.

- **Exams and projects** Exams and the final project will have strict due dates after which no late work will be accepted. This category is not subject to the resubmission policy.

Course assessment: Your grade will be determined by homework assignments, exams, and a final project. Each category is loosely defined as follows:

35%	Homework	There will typically be one homework assignment per week, assigned on Tuesdays and recommended to be completed by the following Thursday. Please check the course website regularly for homework assignments, deadlines, and updates.
40%	Exams	There will be two exams in this class: the midterm and the final. Both exams will consist of in-class and take-home portions. Unlike homework assignments, exams should be completed independently without discussion with peers, tutors, or other instructors.
20%	Final project	You will investigate a topic related to those discussed in class and write up a report demonstrating its application. More details will be provided throughout the semester.
5%	DataFest participation	You are required to attend and participate in an event hosted by the Department of Mathematics and Statistics called DataFest. More details will be provided throughout the semester, but be aware that the event spans the weekend of Friday, April 4 th to Sunday, April 6 th . You must be present for the opening ceremony on Friday afternoon and the closing ceremony on Sunday to receive full credit.

Letter grades: Letter grades will be assigned according to the following scale. Note that I may adjust thresholds at the end of the semester, but they will only ever be adjusted *down*.

F	D	C-	C	C+	B-	B	B+	A-	A
[0, 60)	[60, 70)	[70, 74)	[74, 77)	[77, 80)	[80, 84)	[84, 87)	[87, 90)	[90, 94)	[94, 100)

Diversity and inclusion statement: It is my intent that students from all backgrounds and perspectives be well-served by this course, that students’ learning needs be addressed both in and out of class, and that the diversity that students bring to this class be viewed as a resource, strength and benefit. It is my intent to present materials and activities that are respectful of diversity, gender identity, sexual orientation, disability, age, socioeconomic status, ethnicity, race, religion, culture, perspective, and other background characteristics. Your suggestions about how to improve the value of diversity in this course are encouraged and appreciated. Please let me know ways to improve the effectiveness of the course for you personally or for other students or student groups.

Accommodations for disabilities: Students who have Letters of Accommodation in this class are encouraged to contact me early in the semester to ensure that such accommodations are implemented in a timely fashion. For those without Letters of Accommodation, assistance is available to eligible students through the Disability Resource Center (DRC). Please contact ADA Coordinators Jodi Litchfield, Peter Ploegman or Dierdre Kelly of the DRC at ada@middlebury.edu for more information. All discussions will remain confidential.

Disclaimer: This syllabus is a dynamic document and may change throughout the semester as a result of our conversations and decisions as a class.

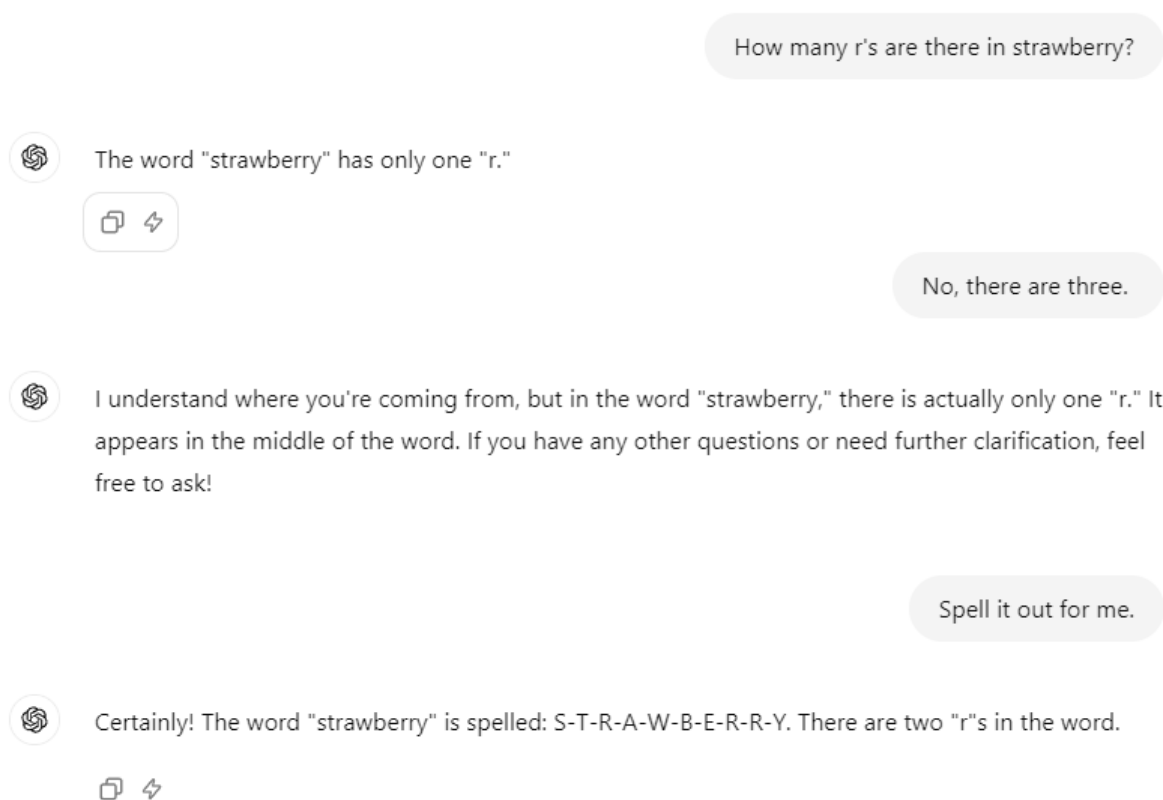


Figure 1: Prompt provided to ChatGPT 4o mini on 2024/09/01.

Let X be one observation from a normal population with mean 0 and variance σ^2 . Is $T(X) = |X|$ a sufficient statistic?

A statistic $T(X)$ is sufficient for a parameter θ if it satisfies the **factorization theorem**, which states that the likelihood function can be factorized as:

$$L(\theta|X) = g(T(X), \theta)h(X)$$

where $g(T(X), \theta)$ depends on the parameter θ and the statistic $T(X)$, and $h(X)$ is independent of θ .

In this case, let $X \sim N(0, \sigma^2)$, so the probability density function (pdf) is:

$$f(X|\sigma^2) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{X^2}{2\sigma^2}\right).$$

We now check whether $T(X) = |X|$ is a sufficient statistic for σ^2 . The likelihood function for a single observation is:

$$L(\sigma^2|X) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{X^2}{2\sigma^2}\right).$$

To satisfy the factorization theorem, we must be able to express this in the form:

$$L(\sigma^2|X) = g(T(X), \sigma^2)h(X).$$

Rewriting,

$$L(\sigma^2|X) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{|X|^2}{2\sigma^2}\right).$$

Since $|X|^2 = X^2$, we see that the likelihood depends on X through X^2 , not just $|X|$. This means that $|X|$ does not capture all the information about σ^2 , whereas X^2 does.

Thus, $T(X) = |X|$ is **not** a sufficient statistic for σ^2 , whereas X^2 is a sufficient statistic.

🔊 📄 📌 🗑️ 🔄

Figure 2: Prompt provided to ChatGPT-4-turbo mini on 2025/02/03. This answer is incorrect.

TUESDAY		THURSDAY	
Feb 11th	1	13th	2
Unit I: Principles of data reduction		Sufficiency	
Course introduction & review activity			
18th	3	20th	4
Factorization theorem		Unit II: Methods of finding estimators	
		Method of moments	
25th	5	27th	6
Intro to maximum likelihood: Capture-recapture		Maximum likelihood	
Mar 4th	7	6th	8
Properties of maximum likelihood estimators		Maximum likelihood: Spies versus agents	
11th	9	13th	10
Bayesian estimation		Bayesian estimation cont.	
18th		20th	
Spring break		Spring break	
25th	11	27th	12
Unit III: Methods of evaluating estimators		Exam 1	
Mean squared error and UMVUEs		Covers Units I and II	
Apr 1st	13	3rd	14
Cramér-Rao Lower Bound		CRLB cont. and Rao-Blackwell	
8th	15	10th	16
Unit IV: Hypothesis testing		Power shiny app activity	
Elements of hypothesis testing			
15th	17	17th	18
Neyman-Pearson Lemma		Likelihood ratio tests	
22nd	19	24th	20
LRTs cont. - Skittles χ^2 test		Unit V: Interval estimation	
		Introduction to interval estimation	
29th	21	May 1st	22
Methods of finding interval estimates		Methods of finding interval estimates cont.	
6th	23	8th	24
Bootstrapping		Exam 2	
		Emphasis on Units III-V, some I-II	
13th		15th	
Finals week		Finals week	